

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant:	Bruce Schofield	Examiner:	John Pauls
Serial No.:	10/805,993	Art Unit:	3686
Conf. No.:	1210		
Filed:	March 22, 2004	Attorney Docket No.:	909430-US-NP
Title:	METHOD AND APPARATUS FOR PROVIDING NETWORK BASED LOAD BALANCING OF MEDICAL IMAGE DATA		

CERTIFICATE OF ELECTRONIC TRANSMISSION

I hereby certify that this document, along with any other papers referred to as being attached or enclosed, is being filed electronically on September 21, 2010.

/John C. Gorecki/
John C. Gorecki, Reg. No. 38,471

M.S. Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

REPLY BRIEF

This reply brief is being filed in response to the Examiner's answer dated July 21, 2010. This reply brief is being filed to further narrow issues to be decided on appeal.

Independent claims 37 and 45 were rejected under 35 USC 103 as unpatentable over Rothschild¹ in view of Primak², and Martin³. In this application, both independent claims recite that the level of complexity of the task to be performed should be used during the load balancing process.^{4 5} The Examiner erred in rejecting these claims since none of the references teaches or

¹ U.S. Patent Application Publication No. 2002/0016718

² U.S. Patent No. 6,389,448

³ U.S. Patent No. 6,263,368

⁴ In claim 37, the method includes the steps of "determining, by the network service, a level of complexity of the task to be performed from the instructions associated with the task" and "selecting ... one of the plurality of image archive resources to be used to perform the task ... using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed."

⁵ In claim 45, the apparatus includes a network service configured to "determine a level of complexity of the task to be performed..." and "select one of the plurality of image archive resources to be used to perform the task ... using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed."

suggests that the level of complexity of the task to be performed should be used to select a processor during the load balancing process.

On page 7 of the Examiner's Answer, the Examiner cites Martin at col. 1, lines 41-44, Col. 3, lines 30-48, Col. 4, lines 3-5 and Col. 8, lines 38-43 as teaching this aspect. At Col. 1, lines 41-44, Martin teaches that tasks should be distributed equally among the individual server computers to balance the overall loading. At Col. 3, lines 30-48 and Col. 4, lines 3-5, Martin teaches that load balancing should be based on network link loading rather than server loading. At Col. 8, lines 38-43, Martin teaches that the dispatch controller looks at a traffic loading table to allocate individual in-bound requests. All of these sections focus on the network link loading, i.e. the number of messages being transmitted by the processors in connection with processing previous tasks. The amount of traffic associated with other previously allocated tasks has nothing to do with determining the complexity of the task that has not yet been allocated.

The Examiner focuses on the fact that Martin mentions network loading, processor loading, and packet or byte counts (Examiner Answer at page 7, line 21- Page 8, line 2) and states that this is the same as the complexity mentioned by applicant in the specification. What applicants actually state, in paragraph 22, is:

The network service identifies the task within the DICOM message and extracts it. The network service determines the level of complexity of the received task, or assigns a priority the task using predetermined criteria (206). Typically, the level of complexity of a task will be based on the amount of resources that are required to execute the task such as processor time or memory, or the complexity may be a function of the time that is required to execute the command.

As is clear from this paragraph, the number of packets or byte counts is not mentioned in paragraph 22 as being related to the complexity of the task that is to be load balanced. Indeed, the network loading, processor loading, packet and byte counts are all quantities that are separate from the complexity of the task that is to be load balanced. Network loading looks at how much message traffic is being to/from individual servers. (Martin at Col. 3, lines 42-46). Processor loading looks at how busy the processor is. (Martin at Col. 2, lines 30-33). Martin counts the number of packets or bytes to determine how busy the network links are; this is part of the network load monitoring. (Martin at Col. 3, line 62 to Col. 4, line 5). All of these quantities tell the load balancer how the system or network is running based on previously allocated tasks, and have nothing to do with determining the complexity of a task that is yet to be allocated.

As noted above, independent claim 37 recites a method that includes the step of “determining, by the network service, a level of complexity of the task to be performed from the instructions associated with the task.” Likewise, independent claim 45 recites an apparatus that will “determine a level of complexity of the task to be performed from the instructions associated with the task.” The combination of Rothschild/Primak/Martin does not teach using the complexity of the task to be performed in connection with load balancing.

Further, the claims both recite that the “complexity” that is used in the load balancing determination is determined from the instructions associated with the task. On page 8 of the Examiner’s Answer, the Examiner states that the combination of Rothschild/Primak/Martin “may or may not specifically disclose the following limitations: *extracting, by the network service, the instructions associated with the task from the medical image data.*” Thus, the Examiner admitted that the cited combination of references does not use instructions associated with the task to be implemented to determine the complexity of the task to be load balanced. The complexity cannot be determined based on the “instructions” unless the instructions are first extracted. However, the Examiner finds that it is well known that the DICOM format includes these instructions and, hence, that it would have been obvious to look at the instructions when load balancing.

Applicants respectfully disagree with this contention. Specifically, as explained in greater detail in the Appeal Brief, none of the three references looks at the complexity of the task that is being load balanced. Accordingly, the fact that the DICOM format includes instructions which specify what the end device is required to do in connection with the images would not have suggested that the complexity of the task should be taken into consideration in the first instance. Rather, the well known DICOM format should be viewed simply as another piece of prior art. The DICOM format does not suggest that the complexity of the task to be implemented should be considered when load balancing any more than any of the other references. More succinctly, the combination of Rothschild/Primak/Martin/DICOM does not suggest looking at the complexity of the task any more than the combination of Rothschild/Primak/Martin. Thus, applicants respectfully submit that the Examiner failed to establish, *prima facie*, that the claims are obvious since none of the references teach or suggest looking at the complexity of the task when implementing load balancing. Further, the combination does not teach or suggest how to assess the complexity. Applicants proposed to extract the instructions associated with the task

from the medical data. The mere fact that the DICOM standard includes this information does not mean that it would have been obvious to use this information for a different purpose – namely for load balancing. Accordingly, applicants respectfully submit that the Examiner erred in rejecting the claims under 35 USC 103.

Response to Arguments Section

The Examiner made several points in the “Response to Arguments” section of the Examiner’s Answer (Examiner’s Answer at pages 14-15) which Applicant will address. Specifically, the Examiner has taken the position that both Primak and Martin teach load balancing based on complexity of the task.

In Primak, when a SYN packet is received from a client computer, a SYN value is randomly assigned to the packet (Primak at Col. 4, lines 40-46). The number assigned to the packet is calculated using a number generation function characterized by an even probability distribution over a fixed range of values. (Primak at Col. 3, lines 60-63). For example, assume that the range of values is 1-10,000. When a SYN packet is received, it is assigned a value from that range but the value can be any number and there is equal probability of having any one of the values assigned to the SYN packet.

The servers, in Primak, are designed to accept responsibility for any SYN packet within their defined range. (Primak at Col. 4, lines 42-46). In the hypothetical example, if there are three servers the first server may be responsible for processing SYN packets having an assigned SYN number in the range 1-3000, the second server may be responsible for SYN values 3001-5000, and the third server may be responsible for SYN values 5001-10000. The ranges of a server’s responsibility are dynamically adjusted based on how busy the server is. (Col. 4, line 63 to Col. 5, line 20). When a SYN packet is received, a pseudorandom number is generated and assigned to the SYN packet. Any server that is responsible for a range of SYN numbers that encompasses the SYN number assigned to the SYN packet will assume responsibility for the SYN packet. Thus, assuming a SYN value of 3098 was assigned to the SYN packet, the second server in this hypothetical would assume responsibility for the SYN packet. The complexity of a new connection, i.e. complexity of a SYN packet (how much processing will be required to process a newly received SYN) is never considered when determining which processor should handle the SYN packet.

On page 14, lines 8-12, the Examiner states that Primak discloses that the SYN packet is used to calculate a connection value for the requested task. This is at least partially correct. Primak teaches that the pseudo-random function mentioned above (referenced by Primak at Col. 3, lines 57-59) preferably uses a portion of the SYN packet as input to the function when generating the pseudo-random number. This does not have anything to do with assessing the complexity of the task to be implemented, however. Rather, Primak uses the packet as partial input to the algorithm that is being used to generate a random number to be assigned to the packet. Assigning a random number to a packet is pretty much the antithesis of assessing and assigning a complexity value to the packet.

Further, the Examiner focuses on the fact that only connection values that fall within the servers' available capacity are accepted. The server's available capacity has nothing to do with the complexity of the task that is to be assigned. Primak does not look at the complexity of the SYN packet. Looking at the available capacity of the server does not provide any information about the complexity of the task that has not yet been assigned to one of the servers.

On page 14, line 13- page 15, line 2, the Examiner focuses on Martin to explain how Martin shows using the complexity of the task in connection with load balancing.

Martin recognizes that assigning tasks based on a strict round robin basis to a set of computers is not ideal, because some requests can result in significantly different processing requirements. (Martin at Col. 2, lines 14-29).

However, Martin's solution is not to look at the complexity of a new task when determining how to load balance that new task. Rather, Martin teaches that the server load should be evaluated and that the load balancing process should avoid sending new tasks to servers that have high usage levels. (Martin at Col. 2, lines 30-48). Martin posits that processor load may not be the only limiting factor (Martin at Col. 2, lines 56-61) and that the network link loading to the server may also be a limiting factor. (Martin at col. 2, line 61 – Col. 3, line 1). Hence, Martin proposes that network loading should be used to make load balancing decisions.

The Examiner stated that Martin concludes that monitoring network link traffic represents the complexity of the requested tasks. (Examiner's Answer at Page 14, lines 18-19). This doesn't make logical sense. Link traffic is generated by the processors when the processors are working on tasks. Thus, the link traffic is associated with tasks that have already been load balanced. Link traffic does not provide an indication of the complexity of a new task that has not

yet been load balanced. While the link traffic may tell the load balancing module how busy the servers are, it does not look at the complexity of the new task. Thus, although Martin acknowledges that different tasks may be of different complexities, Martin doesn't propose to try to figure out in advance how complex a new task will be. Note that the claims require the complexity to be determined and then, the complexity is used to during the step of selecting. Hence, the claims are focused on allocating a new task to one of the servers and are not addressing a situation where the complexity of previously allocated tasks is being assessed. Martin looks at network loading which, like processor loading, tells the load balancing system how busy servers are at executing previously allocated tasks. Network loading, like processor loading, does not look at the complexity of the task that is yet to be assigned when selecting a processor to operate on that task.

Conclusion

Applicants respectfully submit that the Examiner has committed legal error in rejecting the claims of this application, since none of the references look at the complexity of the task to be load balanced when making a load balancing decision. Accordingly the rejection under 35 USC 103 should be reversed.

If any fees are due in connection with this filing, the Commissioner is hereby authorized to charge payment of the fees associated with this communication or credit any overpayment to Deposit Account No. 502246 (Ref: 909430).

Respectfully Submitted

Dated: September 21, 2010

Anderson Gorecki & Manaras LLP
P.O. Box 553
Carlisle, MA 01741
Tel: (978) 264-4001
Fax: (978) 264-9119

/John C. Gorecki/
John C. Gorecki, Reg. No. 38,471